

IN THE DRAWINGS

Please replace Fig. 6 with the attached replacement Fig. 6.

REMARKS

Claims 1-27 were pending. Claims 1, 3, 4, 7, 10, 12, 13, 17-24 and 26 have been amended for clarification purposes. Accordingly, claims 1-27 remain pending subsequent this amendment.

Fig. 6 has been amended to correct an error and bring the figure into conformity with the description. In particular, the flow out of block 618 has been amended to flow to block 620, rather than to block 622. A replacement drawing sheet is attached.

In the present Office Action, claims 17-23 stand rejected under 35 U.S.C. § 101 as being directed to non-statutory subject matter. Applicant has amended each of claims 17-23 and believes these rejections are overcome.

Claims 3-4, 7-9, 12-13, 16, 19, 20, 23 and 26 stand rejected under 35 U.S.C. § 112. In response, Applicant has amended claims 3, 12, 19 and 26 to clarify the nature of “mounting” and “unmounting”. With these amendments, and the accompanying description, Applicant believes anyone skilled in the art will comprehend these terms and the rejections are overcome.

In the present Office Action, claims 1, 6, 10, 15, 17, 22, 24 and 27 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,567,865 (hereinafter “Araki”). Claims 1, 5, 6, 10, 14, 15, 17, 21, 22, 24, and 27 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,044,367 (hereinafter “Wolff”). Finally, claims 1, 2, 6, 10, 11, 15, 17, 18, 22, 24, 25, and 27 stand rejected under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,675,268 [erroneously cited as U.S. Patent No. 6,657,268 in paragraph 6] (hereinafter “DeKoning”). Applicant respectfully traverses these rejections and requests reconsideration of all claims in view of the following discussion.

Araki is generally directed to a method for controlling process accesses to storage on an extent basis. With respect to prior art, Araki notes that prior systems may lock access on a volume basis, file, or sub-volume basis, thereby preventing access to the locked volume, file, or sub-volume by other processes. As noted by Araki, such exclusive access may be unnecessary when the processes are not seeking access to the same data. Therefore, Araki discloses determining whether accesses overlap – on an extent basis – in order to determine exclusivity. For example, Araki teaches the following:

“In view of the above and other objects which will become apparent as the description proceeds, there is provided according to an aspect of the present invention a storage system which is comprised of a control unit incorporating a control memory, wherein information concerning the extent (range) of an input/output processing request (hereinafter also referred to as the input/output processing request extent) which is transferred from a given one of plural host processors to the control unit upon issuance of the input/output processing request from the former is stored in the control memory with a view to realizing the exclusive control for a plurality of input/output processing requests issued from a plurality of host processors to one logical device by making use of the extent information mentioned above. When an input/output processing request is newly issued, decision is made as to whether or not any input/output processing request whose extent overlaps with that of the newly issued input/output processing request is being executed. Unless overlap of the extents is found, the input/output processing as requested currently is executed continuously. On the other hand, when the extent overlap is found, the newly issued input/output processing request is stored in the control memory as the input/output processing request waiting for execution.” (col. 2, lines 7-30).

“The extent information 401 contains a start address and an end address of the relevant logical device determined on the basis of the extent of the input/output processing request designated by the command "Define Extent", to thereby define the extent of the logical device over which the input/output processing is to be performed. The control unit of the storage system according to the present invention is so designed or programmed as to perform the exclusive control on the basis of the extent information 401. Consequently, when the extent information 401 of a given input/output processing overlaps with that of the other input/output processing which is being currently executed, the given input/output processing is forced to wait for the start of

execution thereof, which is referred to as occurrence of wait event. In this case, information concerning the relevant wait-event occurrence time 402 is stored in the input/output processing control table.” (col. 5, lines 51-67).

Accordingly, Araki discloses controlling access to a shared resource by processes based upon whether or not the accesses are directed to overlapping extents. Consequently, Araki is directed to a different problem and does not disclose the features of the independent claims. For example, Araki nowhere discloses “initiating a storage re-allocation procedure in said computer network, wherein said re-allocation procedure is configured to re-allocate a first storage from a first host in said computer network to a second host in said computer network.” Such features are wholly absent from Araki. Further, the remaining features of claim 1 “determining whether I/O corresponding to said first storage is in progress, in response to detecting said re-allocation procedure has been initiated; and halting said re-allocation procedure in response to determining I/O corresponding to said first storage is in progress.” (emphasis added). Col. 4, lines 23-38 of Araki, cited in paragraph 4 of the Office Action, merely states that a controller either returns data from a cache, if stored therein, or the targeted device in response to a read request. Therefore, Applicant submits claim 1 (and claims 10, 17 and 24 for similar reasons) is patentably distinguishable from Araki.

Also cited as anticipating the claimed invention is Wolff. Wolff is generally directed to increasing throughput in a client/server system by dividing requests into administrative and data portions. Each of the portions may then be concurrently handled by different servers. However, Wolff merely discloses a server receiving and queuing I/O requests directed to a particular network resource. For example, Wolff discloses:

“Computer networks require file servers which frequently operate under the client/server paradigm. Under this paradigm, multiple clients make I/O requests which are directed to a particular resource on the network. A server on the network receives and carries out the I/O requests. When a server receives multiple I/O requests, the server queues them and then services them one at a time. Once a queue begins to accumulate, subsequent I/O requests must sit in the queue

until the previous I/O requests are serviced. As a result, the server can become a bottleneck in the network.” (Wolff, col. 1, lines 56-65).

Accordingly, like Araki, Wolff does not disclose “initiating a storage re-allocation procedure in said computer network, wherein said re-allocation procedure is configured to re-allocate a first storage from a first host in said computer network to a second host in said computer network.” Further, Wolff does not disclose “determining whether I/O corresponding to said first storage is in progress, in response to detecting said re-allocation procedure has been initiated; and halting said re-allocation procedure in response to determining I/O corresponding to said first storage is in progress.”

Finally, Koning is generally directed to preventing thrashing between controllers in a network. To that end, Koning discloses using a timing mechanism which provides ownership to a controller for a predetermined period of time – irrespective of whether or not the controller is currently accessing the owned device. In addition, Koning discloses taking into consideration the topology of the network and the priority of a request when making a decision regarding transfer of ownership. For example, Koning teaches:

“The present invention prevents excessive thrashing between array controllers in a storage area network (SAN) when more than one host device is attempting to access the same logical data volume through different array controllers. Thrashing is prevented by granting ownership of a logical data volume to a particular array controller for a period of exclusivity. During the exclusivity period, ownership of the logical data volume may not be transferred to another array controller. The logical data volume is said to be "sticky" to the particular array controller for the exclusivity, or "sticky," time period, e.g. about five to ten seconds or longer depending on the situation.

The present invention also considers whether a host device may lose access to a needed logical data volume that it is currently accessing before permitting transfer of ownership of the logical data volume from the current-owner array controller to a nonowner array controller that has received a data access request to the logical data volume. Therefore, before transferring ownership of the logical data volume, the nonowner array controller examines a portion of the current topology of the SAN to determine whether any host device that requires access to the logical data volume would lose access to the

logical data volume if ownership were to be transferred. The topology is the set of all connections between the host devices and the storage arrays. In this case, the nonowner array controller determines those available connections to the storage array of which it is a part, including connections through other array controllers in the same storage array. If, according to the available connections, any host device would lose access to the logical data volume upon transferring ownership from the current-owner array controller to the nonowner array controller, then preference is given to not transferring ownership. Ownership transfer is granted only if the requesting host device cannot establish a data transfer path to the current-owner array controller and the requesting host device has priority for accessing the logical data volume over the host device that is currently accessing the logical data volume.” (col. 2, lines 30-67).

“Prior art array controllers permit the occurrence of the AVT process as soon as the data access request is received by the non-owner array controller, and prior art RDAC software supports this occurrence. The present invention, however, can reject the ownership transfer, or deny the data access request and refuse to permit the AVT process to occur, in certain circumstances. In that case, the array controller 142 sends an appropriate error message to the RDAC software 140 in the host device 104 describing the reason for not processing the data access request. For example, ownership transfer of the logical data volume 128 may be rejected when an insufficient amount of time has passed since the last AVT process transferred ownership of the logical data volume 128 between the array controllers 142 and 144. Additionally, ownership transfer can be rejected when such transfer would cause another host device 106 and/or 108 to lose access to the logical data volume 128 when the other host device 106 and/or 108 requires such access.” (col. 7, lines 25-43).

Accordingly, Koning is directed to reducing thrashing by implementing a minimum period of time which must pass before ownership may be transferred. In addition, the topology of the network may be examined when making a transfer decision. However, Applicant submits Koning does not disclose or suggest all of the features of the independent claims. Koning does not disclose “initiating a storage re-allocation procedure in said computer network, wherein said re-allocation procedure is configured to re-allocate a first storage from a first host in said computer network to a second host in said computer network.” The mere transfer of ownership between controllers, which may

occur regularly and repeatedly, disclosed in Koning is not equivalent to re-allocating storage as claimed and described. Further, Koning does not disclose or suggest “determining whether I/O corresponding to said first storage is in progress, in response to detecting said re-allocation procedure has been initiated; and halting said re-allocation procedure in response to determining I/O corresponding to said first storage is in progress.” Rather, Koning makes a transfer decision based upon a determination as to whether a sufficient period of time has passed and/or whether or not the network topology indicates a host may lose access in the event of a transfer. Accordingly, Applicant submits the presently claimed invention is distinguishable from Wolff as well.

In view of the above, Applicant believes each of the independent claims are patentably distinguishable from the cited art. Accordingly, each of the dependent claims are patentable as well. In addition, the dependent claims recite additional features which are neither taught nor suggested by the cited art. For example, claim 3 recites the additional features “unmounting said first host from said first storage and configuring said first host to bypass mounting said first storage upon a subsequent reboot; and completing said re-allocation procedure; wherein mounting said first storage comprises enabling access to the first storage, and wherein unmounting said first storage comprises disabling access to the first storage.” As each of the cited references is directed to a different problem, the features of claim 3 are wholly absent from the cited art. Further, the additional features recited in claim 4, 7, 8 and 9 are nowhere found in the cited art.

CONCLUSION

Applicant submits the application is in condition for allowance, and an early notice to that effect is requested.

If any extensions of time (under 37 C.F.R. § 1.136) are necessary to prevent the above referenced application(s) from becoming abandoned, Applicant(s) hereby petition for such extensions. If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5181-72700/RDR.

Also enclosed herewith are the following items:

- ☒ Return Receipt Postcard
- ☒ Replacement Drawing (Fig. 6)

Respectfully submitted,



Rory D. Rankin
Reg. No. 47,884
ATTORNEY FOR APPLICANT(S)

Meyertons, Hood, Kivlin,
Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: 2/1/05